arXiv:1604.04052v1 [math.NA] 14 Apr 2016

MEMORY-EFFICIENT RECYCLING OF LARGE KRYLOV-SUBSPACES FOR SEQUENCES OF HERMITIAN LINEAR SYSTEMS

MARTIN P. NEUENHOFEN* AND SVEN GROß†

Abstract. We present a new short-recurrence residual-optimal Krylov subspace recycling [5] method for sequences of Hermitian systems of linear equations with a fixed system matrix and changing right-hand sides. Such sequences of linear systems occur while solving, e.g., discretized time-dependent partial differential equations.

With this new method it is possible to recycle large-dimensional Krylov-subspaces with smaller computational overhead and storage requirements compared to current Krylov subspace recycling methods as e.g. R-MINRES [1].

In this paper we derive the method from the residual-optimal preconditioned conjugate residual method (PCR, [27, p. 182]) and discuss implementation issues. Numerical experiments illustrate the efficiency of our method.

AMS subject classifications. 65F10, 65N22, 76M10, 93A15, 93C05, 65F50.

Key words. Sequence of linear systems, iterative linear solvers, Krylov subspace recycling, Krylov methods, R-MINRES, MINRES, CR, PCR.

1. Introduction.

We present a new Krylov-subspace (KSS) method for iterative numerical solution of sequences of n_{Eqns} Hermitian linear systems

$$\mathbf{A}\mathbf{x}^{(\iota)} = \mathbf{b}^{(\iota)}, \qquad \iota = 1, \dots, n_{\text{Eqns}}\,, \tag{1.1}$$

with fixed $\mathbf{A} \in \mathbb{C}^{N\times N}$, where the right-hand sides (RHS) $\mathbf{b}^{(\iota+1)} \in \mathbb{C}^N$ depend on the former solution $\mathbf{x}^{(\iota)}$ for the RHS $\mathbf{b}^{(\iota)}$. Thus the systems must be solved *one after the other*. Such situations occur e.g. when applying an implicit time stepping scheme to numerically solve a non-stationary PDE. Areas of application are e.g. structural dynamics [5], topology optimization [1], circuit analysis [8], fluid dynamics [7] and model reduction [4].

For simplicity, in this paper we present the case where $\mathbf{A}$ is regular, but from [20, 21] it is well-known that CR- and MINRES-type methods can also be applied to singular systems.

1.1. Krylov-Subspace Methods.

Using a similar notation as in [5], projection methods build a search space $\mathcal{U} \subset \mathbb{C}^N$ with image $\mathcal{C} = \mathbf{A}\mathcal{U}$, in which an approximate solution $\mathbf{x}$ for *one* single system is searched, such that the *residual* $\mathbf{r} = \mathbf{b} - \mathbf{A}\mathbf{x}$ is orthogonal to a test space $\mathcal{P} \subset \mathbb{C}^N$. For KSS methods

$$\mathcal{U} = \mathbf{M}^{-1} \cdot \mathcal{K}_m(\mathbf{A}\mathbf{M}^{-1}; \mathbf{b}) = \mathbf{M}^{-1} \cdot \operatorname*{span}_{i=0,\dots,m-1} \{(\mathbf{A}\mathbf{M}^{-1})^i\, \mathbf{b}\} \tag{1.2}$$

is chosen[1], where $\mathbf{M}^{-1}$ is a cheap approximation to the inverse of $\mathbf{A}$, i.e. a preconditioner (or identity when no preconditioning is used). $\mathcal{K}_m$ is called *Krylov-subspace* of dimension m. Coming from linear iterative methods or iterative refinement, this space is the most natural choice, as it adds (in case of using a reasonable preconditioner)

*Martin.Peter.Neuenhofen@rwth-aachen.de, RWTH Aachen, Germany

†gross@igpm.rwth-aachen.de, Institut für Geometrie und Praktische Mathematik (IGPM), RWTH Aachen, Germany.

[1]For an easier presentation we assume $\mathbf{x} = 0$ as initial guess in the introduction.

the correction guess from the current residual to the search space for the solution, i.e. we update

$$\mathcal{U} := \mathcal{U} + \operatorname{span}\{\mathbf{M}^{-1}\mathbf{r}\}\,. \tag{1.3}$$

KSS methods augment the search space by one dimension in each iteration. For each additional dimension at least one matrix-vector-product (MV) with $\mathbf{M}^{-1}$ and $\mathbf{A}$ has to be computed, which for short recurrence methods dominate the computation time of the whole solution process. Internally KSS methods build basis matrices $\mathbf{U}, \mathbf{C}, \mathbf{P} \in \mathbb{C}^{N\times m}$ for the three spaces.

Two big goals for KSS methods are 1. short recursions such that only a limited number of columns of the basis matrices must be stored (necessary due to hardware limitations) and also leading to lower computational costs per iteration (efficiency); 2. a good test space $\mathcal{P}$ that yields small residual norms to ensure a good approximate solution w.r.t. the search space (important for convergence). For Hermitian systems preconditioned CR and MINRES satisfy both conditions.

For more information about KSS methods we refer to [27].

1.2. Krylov-Subspace Recycling. When it comes to solving sequences of linear systems as in (1.1), it seems desirable to maintain the construction idea from (1.3) over the whole sequence to possibly save some of the MV-computations. I.e., *one* high dimensional search space is successively constructed for all (or many) RHS-es of the sequence, where always the current residual (of the current RHS) is added to the search space. Such a construction leads to sums of KSS-es, because KSS information from former solution processes is reused for construction of the search space to solve the current system. Consequently such methods are called *Krylov subspace recycling methods* or short *recycling methods* [7, 5, 1].

In [4] a GCR-based (i.e. $\mathbf{P} = \mathbf{C}$, $\mathbf{C}$ orthogonal) KSS recycling method is constructed by simply adding the current (preconditioned) residual as new search direction to the search space basis matrix $\mathbf{U}$, cf. [4, p. 5] (in their notation $\mathbf{p}_k$ are the columns of $\mathbf{U}$), [16, Algo. 1]. This method can be characterized as a *full* recycling method, as it stores the full information from all former systems. However the authors note that this method needs modifications due to memory limits and discuss remedies.

In case that the basis matrices $\mathbf{U}, \mathbf{C}$ have become too large to be completely stored in memory, there are two remedies known from the literature: *short representations* [16] and basis reduction techniques (e.g. deflation [15], optimal truncation [14]), also summarized under the term *compressings* in [8].

The idea of the first approach (short representations) is to keep all the subspace information from the bases $\mathbf{U}, \mathbf{C}$ — but using less storage — by exploiting the structure of Krylov subspaces. This is what we also propose in the present paper. To our best knowledge, besides [16], there is no other literature on this short representation technique.

In the other approach, based on compressings, the idea is to reduce the basis dimension by a reduction rule of the form

$$\begin{aligned} \mathbf{U}_c &:= \mathbf{U}\,\boldsymbol{\eta} \\ \mathbf{C}_c &:= \mathbf{C}\,\boldsymbol{\eta}\,, \end{aligned} \tag{1.4}$$

where $\boldsymbol{\eta}$ has fewer columns than $\mathbf{U}$ and $\mathbf{C}$, leading to compressed basis matrices $\mathbf{U}_c, \mathbf{C}_c$. There are different approaches for the reduction, where the most common

choices are to use Ritz vectors (called *deflation*) [5, 15] eliminating eigenvalues that slow down convergence, and *optimal truncation* [14].

In general in compressing methods some information is lost. So by construction the efficiency of this approach is limited compared to that of the short representation approach, as the latter is capable of conserving *all* subspace information computed so far. This property only holds for the case of a constant system matrix as considered in (1.1).

In the following we review compressing-based KSS recycling methods from the literature.

Actually KSS recycling methods originated from restart strategies for GMRES(k,m) and GCRO(k,m), cf. [10, 11, 14]. For both methods after m iterations (of the respective KSS method) a restart is performed. During this restart a k-dimensional search space basis matrix $\mathbf{U}_c \in \mathbb{C}^{N\times k}$ $(k < m)$, obtained from the iterations before by (1.4), is conserved and extended by a matrix $\mathbf{W} \in \mathbb{C}^{N\times m}$ either for GMRES(k,m) by *augmentation*

$$\mathbf{A}[\mathbf{U}_c, \mathbf{W}] = [\mathbf{C}_c, \overline{\mathbf{W}}] \begin{bmatrix} \mathbf{I} & \mathbf{B} \\ \mathbf{0} & \overline{\mathbf{H}} \end{bmatrix},$$

or for GCRO(k,m) by *orthogonalization*, i.e.

$$(\mathbf{I} - \mathbf{C}_c\mathbf{C}_c^H)\,\mathbf{A}\,\underline{(\mathbf{I} - \mathbf{C}_c\mathbf{C}_c^H)^H}\mathbf{W} = \overline{\mathbf{W}}\,\overline{\mathbf{H}}. \tag{1.5}$$

In the equations above we have $\mathbf{W} = [\mathbf{w}_1, \ldots, \mathbf{w}_m], \overline{\mathbf{W}} = [\mathbf{W}, \mathbf{w}_{m+1}] \in \mathbb{C}^{N\times(m+1)}$, $\overline{\mathbf{H}} \in \mathbb{C}^{(m+1)\times m}$ upper Hessenberg, representing the underlying Arnoldi process, and $[\mathbf{C}_c, \overline{\mathbf{W}}]$ *orthogonal*, i.e. $[\mathbf{C}_c, \overline{\mathbf{W}}]^H[\mathbf{C}_c, \overline{\mathbf{W}}] = \mathbf{I}$, with $\mathbf{I}$ the identity.

When not using these techniques for a restart but for recycling, the matrix $\mathbf{U}_c$ in this context is called *recycling basis* [5].

In (1.5) the underlined factor $(\mathbf{I} - \mathbf{C}_c\mathbf{C}_c^H)^H$ can be dropped due to $\mathbf{C}_c^H\mathbf{W} = \mathbf{0}$ but is written here to emphasize the symmetry of the modified system matrix. For R-MINRES [1] this symmetry is exploited to orthogonalize $\mathbf{W}$ by short recursions. As GMRES, MINRES, GCR, CR, CG, BiCG and BiCR can be all derived from a common principle there are multiple publications on each of these KSS variants for different preconditioning and compressing strategies, e.g. for BiCG [2], BiCGstab [3] and GMRES [12, 13].

1.3. Motivation and Outline. In the following we propose a KSS recycling method for sequences of systems (1.1), that is capable of

1. recycling a high dimensional search space $\mathcal{U}$, built up by rule (1.3);
2. finding a residual-optimal[2] solution in that space;
3. and afterwards improving the solution by iteratively extending $\mathcal{U}$ by rule (1.3) only with short recurrences, and computing the residual-optimal solution in the whole space (also only using short recurrences).

To find a residual-optimal solution in an m-dimensional recycled search space $\mathcal{U}$ of a former system, the computational costs are as follows. Let $k, J \in \mathbb{N}$ with $m = kJ$ (the meaning and choice of k, J will be discussed later). Our method needs $\mathcal{O}(J)$ MV-s with $\mathbf{A}$ and storage for $\mathcal{O}(k)$ column vectors of length N (instead of m column

[2] In case of preconditioning $\|\mathbf{M}^{-1}\mathbf{r}\|_{\mathbf{M}}$ is minimized. For details we refer to section 2.1.3 .

vectors for a naive full recycling strategy). Then post-iterations are carried out where $\mathcal{U}$ is extended by rule (1.3) by short recurrences and the residual-optimal solution to the current RHS is found in this extended search space, which requires 1 MV per post-iteration.

We call the method SR-PCR-ap, which stands for *Short Representation based Recycling for Preconditioned Conjugate Residual with a-posteriori Optimality.*

In section 2 we explain the general ideas of our method and sketch the implementation. Section 3 shows numerical experiments, which illustrate the ability of our method to recycle large KSS-es with small memory consumption. Finally, section 4 gives a conclusion and an outlook.

2. SR-PCR-ap. Our method is based on Preconditioned Conjugate Residual. To clarify notation and properties, we first review this method, as it is applied to solve a single system. Then we recall its properties and introduce our recycling idea.

2.1. Preconditioned Conjugate Residual.

2.1.1. Notation. Throughout this text $\mathbf{U}, \mathbf{D}, \mathbf{V} \in \mathbb{C}^{N\times m}$ denote basis matrices with column vectors $\mathbf{u}_i, \mathbf{d}_i, \mathbf{v}_i \in \mathbb{C}^N$, $i = 1, \ldots, m$. $\overline{\mathbf{U}}, \overline{\mathbf{D}}, \overline{\mathbf{V}}$ consist of the same basis matrices with one additional column to the right. By $\mathbf{U}_{:,1:J:m}$ we denote the basis matrix that consists only of each J^{th} column of $\mathbf{U}$, starting with the first: $\mathbf{U}_{:,1:J:m} = [\mathbf{u}_1, \mathbf{u}_{1+J}, \mathbf{u}_{1+2J}, \ldots] \in \mathbb{C}^{N\times\lfloor m/J\rfloor}$.

Let $\mathbf{e}_i$ denote the ith canonical unit vector, $\mathbf{1}$ a vector that contains only ones, $\mathbf{0}$ a zero vector and $\mathbf{O}$ a zero matrix. We use the symbol $\mathbf{x}_0$ as an *initial guess* for a solution and $\mathbf{r}_0 := \mathbf{b} - \mathbf{A}\mathbf{x}_0$ as *initial residual*; for a sequence of RHS-es we augment this notation by indices, e.g. $\mathbf{x}_0^{(\iota)}, \mathbf{r}_0^{(\iota)}$ for $\iota = 1, 2, \ldots$.

We write $\mathbf{T} \in \mathbb{C}^{m\times m}$ for a Hermitian tridiagonal matrix. $\overline{\mathbf{T}}$ consists of $\mathbf{T}$ with an additional row at the bottom. Its entries are denoted by

$$\overline{\mathbf{T}} = \begin{bmatrix} \alpha_1 & \beta_2 & & \\ \beta_2 & \ddots & \ddots & \\ & \ddots & \ddots & \beta_m \\ & & \beta_m & \alpha_m \\ & & & \beta_{m+1} \end{bmatrix} \in \mathbb{C}^{(m+1)\times m}.$$

2.1.2. The Conjugate Residual Method (CR). Consider an orthogonal sequence $\mathbf{v}_1, \mathbf{v}_2, \ldots \in \mathbb{C}^N$ with $\mathbf{v}_i = \mathbf{A}\mathbf{u}_i$ obtained from a Lanczos iteration $\beta_{i+1}\mathbf{v}_{i+1} = \mathbf{A}\mathbf{v}_i - \alpha_i\mathbf{v}_i - \beta_i\mathbf{v}_{i-1}$ for $i = 1, \ldots, m$ and $\beta_1 \equiv 0$. Equivalently, in matrix notation we obtain

$$\mathbf{A}\mathbf{U} = \mathbf{V}$$
$$\mathbf{A}\mathbf{V} = \overline{\mathbf{V}}\,\overline{\mathbf{T}}$$

with $\overline{\mathbf{V}}$ orthogonal, i.e. $\overline{\mathbf{V}}^H\overline{\mathbf{V}} = \mathbf{I}$. Starting from $\mathbf{u}_1 \parallel \mathbf{r}_0$ the KSS $\mathcal{K}_m(\mathbf{A}; \mathbf{r}_0)$ is spanned by the columns of $\mathbf{U}$. We call $\mathbf{V}$ the image of $\mathbf{U}$. CR constructs a residual-optimal solution in $\mathbf{x}_0 + \mathit{range}(\mathbf{U})$ by subsequently orthogonalizing the residual w.r.t. each column of $\mathbf{V}$ and similarly updating the numerical solution $\mathbf{x}$ with the according column of $\mathbf{U}$.

As $\mathbf{T}$ is tridiagonal, for orthogonality each new column of $\mathbf{V}$ must only be orthogonalized w.r.t. its two left neighboring columns, which makes CR a short recurrence

KSS method. CR is algebraically equivalent to MINRES [19], as both methods compute the residual-optimal solution. Orthogonality and optimality (minimal length of $\mathbf{r}$) are measured w.r.t. the Euclidian scalar-product $\langle\cdot,\cdot\rangle$ and norm $\|\cdot\|$.

2.1.3. The Preconditioned Conjugate Residual Method (PCR). Now we review preconditoning for CR. We write $\mathbf{M} \in \mathbb{C}^{N\times N}$ for a preconditioner with $\mathbf{M} \approx \mathbf{A}$, which must be Hermitian positive definite. Then instead of Euclidian measures we use $\langle\cdot,\cdot\rangle_{\mathbf{M}}$ and $\|\cdot\|_{\mathbf{M}}$, the $\mathbf{M}$-scalar-product and its induced norm[3]. The Lanczos decomposition becomes

$$\begin{aligned} \mathbf{M}^{-1}\mathbf{A}\mathbf{U} &= \mathbf{M}^{-1}\mathbf{D} = \mathbf{V} \\ \mathbf{M}^{-1}\mathbf{A}\mathbf{V} &= \overline{\mathbf{V}}\,\overline{\mathbf{T}} \end{aligned} \tag{2.1}$$

with $\mathbf{D} = \mathbf{A}\mathbf{U}$, where $\overline{\mathbf{V}}$ is then *M-orthogonal*, i.e. $\overline{\mathbf{V}}^H\mathbf{M}\overline{\mathbf{V}} \equiv \overline{\mathbf{D}}^H\overline{\mathbf{V}} = \mathbf{I}$. $\mathbf{u}_1$ is chosen parallel to $\mathbf{M}^{-1}\mathbf{r}_0$.

Now the *preconditioned* CR method (PCR) does not minimize $\|\mathbf{r}\|$ but $\|\mathbf{M}^{-1}\mathbf{r}\|_{\mathbf{M}} \equiv \sqrt{\mathbf{r}^H\mathbf{M}^{-1}\mathbf{r}}$. For readability in the remainder, when speaking about orthogonality and optimality, these properties are always considered with respect to $\langle\cdot,\cdot\rangle_{\mathbf{M}}$ and $\|\cdot\|_{\mathbf{M}}$. For orthogonality w.r.t. $\langle\cdot,\cdot\rangle_{\mathbf{M}}$ we write $\perp_{\mathbf{M}}$. We keep on calling $\mathbf{V}$ *orthogonal* and the *image* of $\mathbf{U}$. If no preconditioner is used, i.e., $\mathbf{M} = \mathbf{I}$, we have $\langle\cdot,\cdot\rangle_{\mathbf{M}} \equiv \langle\cdot,\cdot\rangle$, and $\mathbf{v}_i \equiv \mathbf{d}_i$, $i = 1,\ldots,m+1$.

Algorithm 1 Preconditioned Conjugate Residual variant

procedure PCR($\mathbf{M}, \mathbf{A}, \mathbf{b}, \mathbf{x}_0, m$)
  $\tau := 1, \quad \mathbf{u} := \mathbf{0}, \quad \mathbf{v} := \mathbf{0}$
  $\mathbf{x} := \mathbf{x}_0, \quad \hat{\mathbf{r}} := \mathbf{M}^{-1}(\mathbf{b} - \mathbf{A}\mathbf{x}_0)$ *// $\mathbf{r}_j = \mathbf{b} - \mathbf{A}\mathbf{x}_j$, $j = 0,\ldots,m+1$*
  for $j := 0, 1, 2, \ldots, m$ **do** *// use $\|\hat{\mathbf{r}}\|_2$ or θ as termination criterion*
    $[\hat{\mathbf{v}}, \hat{\mathbf{d}}, \hat{\mathbf{u}}] :=$`A_func`$(\mathbf{M}, \mathbf{A}, \hat{\mathbf{r}})$
      // `A_func` *is:* $\hat{\mathbf{u}} := \hat{\mathbf{r}}$, $\hat{\mathbf{d}} := \mathbf{A}\hat{\mathbf{r}}$, $\hat{\mathbf{v}} := \mathbf{M}^{-1}\hat{\mathbf{d}}$
    $\xi := \langle\hat{\mathbf{d}}, \mathbf{v}\rangle$
    $\mathbf{u} := \hat{\mathbf{u}} - \xi/\tau\,\mathbf{u}, \quad \mathbf{v} := \hat{\mathbf{v}} - \xi/\tau\,\mathbf{v}$
    $\hat{\tau} := \langle\hat{\mathbf{d}}, \mathbf{v}\rangle$
      // columns of $\mathbf{U}$*:* $\mathbf{u}_{j+1} := 1/\sqrt{\hat{\tau}}\,\mathbf{u}$
      // columns of $\mathbf{V}$*:* $\mathbf{v}_{j+1} := 1/\sqrt{\hat{\tau}}\,\mathbf{v}$
      // but columns of $\mathbf{D}$*:* $\mathbf{d}_{j+1} \neq 1/\sqrt{\hat{\tau}}\,\hat{\mathbf{d}}$ *- not required*
      // estimator: $\theta := \hat{\tau} + \xi^2/\tau \equiv \|\mathbf{M}^{-1}\mathbf{A}\mathbf{r}_j\|^2_{\mathbf{M}}$
    $\alpha_j := (\tau - \xi)/\eta, \quad \beta_{j+1} := -\sqrt{\tau\hat{\tau}}/\eta$ *// only if $\overline{\mathbf{T}}$ is of interest*
    $\eta := \langle\hat{\mathbf{d}}, \hat{\mathbf{r}}\rangle, \quad \tau := \hat{\tau}$
    $\mathbf{x} := \mathbf{x} + \eta/\tau\,\mathbf{u}, \quad \hat{\mathbf{r}} := \hat{\mathbf{r}} - \eta/\tau\,\mathbf{v}$ *// $\hat{\mathbf{r}} \equiv \mathbf{M}^{-1}\mathbf{r}_{j+1}$*
  end for
  return $\mathbf{x}$
end procedure

2.1.4. Implementation, Properties and Recycling of PCR. Algorithm 1 gives an implementation of the PCR method as described in section 2.1.3.

[3] i.e. $\langle\mathbf{u},\mathbf{v}\rangle_{\mathbf{M}} := \langle\mathbf{M}\mathbf{u},\mathbf{v}\rangle$, $\|\mathbf{u}\|_{\mathbf{M}} := \sqrt{\langle\mathbf{u},\mathbf{u}\rangle_{\mathbf{M}}}$.

For given $\mathbf{M}, \mathbf{A}, \mathbf{b}, \mathbf{x}_0$, PCR constructs in m iterations the search space

$$\mathcal{U} = \mathbf{M}^{-1} \cdot \mathcal{K}_m(\mathbf{A}\mathbf{M}^{-1}; \mathbf{d}) \tag{2.2}$$

with $\mathbf{d} = \mathbf{r}_0$ and computes a solution $\mathbf{x} \in \mathbf{x}_0 + \mathcal{U}$ such that $\|\mathbf{M}^{-1}\mathbf{r}\|_{\mathbf{M}}$ is minimized, with $\mathbf{r}$ the associated residual to $\mathbf{x}$. This is equivalent to

$$\begin{aligned} \mathbf{x} &= \mathbf{x}_0 + \mathbf{U}\mathbf{V}^H\mathbf{r}_0 \equiv \mathbf{x}_0 + \mathbf{U}\mathbf{U}^H\mathbf{A}\mathbf{M}^{-1}\mathbf{r}_0 \\ \mathbf{r} &= \mathbf{r}_0 - \mathbf{D}\mathbf{V}^H\mathbf{r}_0\,, \end{aligned} \tag{2.3}$$

where $\mathbf{U}$ is a basis of $\mathcal{U}$ and $\mathbf{V}$ its orthogonal image[4].

Eqn. (2.3) also provides the residual-optimal solution $\mathbf{x}$ in the search space $\mathbf{x}_0 + \mathcal{U}$ if we replace $\mathcal{U}$ in (2.2) by choosing a different vector $\mathbf{d}$, e.g. a residual from a former system. Consider as an example the two systems $\mathbf{A}\mathbf{x}^{(1)} = \mathbf{b}^{(1)}$ and $\mathbf{A}\mathbf{x}^{(2)} = \mathbf{b}^{(2)}$. Assume that we have stored the basis matrix $\mathbf{U}$ of $\mathcal{U}$ from the first solution process (with $\mathbf{d} = \mathbf{r}_0^{(1)}$). Then by virtue of formula (2.3) we could compute an optimal solution to the second system in $\mathbf{x}_0^{(2)} + \mathcal{U}$, in which $\mathbf{x}_0^{(2)}$ is the initial guess for the second system with respective residual $\mathbf{r}_0^{(2)}$.

However, in general $\mathbf{U}$ does not fit into the memory. This is why in the next subsection we introduce a scheme to store $\mathbf{U}$ with considerably smaller memory requirements. A discussion on how to address stability issues can be found in section 2.3.

2.2. Basic Idea: Block Krylov Matrices. We see from Algorithm 1 that we can get columns of $\overline{\mathbf{U}}$ and all the entries of $\overline{\mathbf{T}}$ on the fly while solving for one RHS. The idea is now to store *enough* of this data to be able to cheaply compute matrix-vector-products with $\mathbf{U}$ and $\mathbf{U}^H$ when using formula (2.3) to compute residual-optimal solutions for subsequent RHS-es.

The key idea is the notion of a block Krylov matrix as in the following definition.

Definition 1 (Block Krylov Matrix). *Let $\boldsymbol{Z} \in \mathbb{C}^{N\times N}$ and $\boldsymbol{B} \in \mathbb{C}^{N\times q}$ be two arbitrary matrices. For $d \in \mathbb{N}$ with $d\,q \leqslant N$ we define the* block Krylov matrix *$K_d(\boldsymbol{Z}; \boldsymbol{B})$ by*

$$K_d(\boldsymbol{Z}; \boldsymbol{B}) := [\boldsymbol{B}, \boldsymbol{Z}\boldsymbol{B}, \boldsymbol{Z}^2\boldsymbol{B}, \ldots, \boldsymbol{Z}^{d-1}\boldsymbol{B}] \in \mathbb{C}^{N\times(d\,q)}\,.$$

The following lemma enables the computation of matrix-vector-products with $\mathbf{U}$ and $\mathbf{U}^H$ without storing the full matrix.

Lemma 1 (Short Representation). *Let $\boldsymbol{M}, \boldsymbol{A} \in \mathbb{C}^{N\times N}$, $\boldsymbol{U} \in \mathbb{C}^{N\times m}$, $\boldsymbol{T} \in \mathbb{C}^{m\times m}$, as given in* (2.1)*, with $m = kJ$, $k, J \in \mathbb{N}$. Define the permutation matrix $\boldsymbol{\Pi} \in \mathbb{C}^{m\times m}$ column-wise by*

$$\boldsymbol{\Pi}\,\boldsymbol{e}_{1+iJ+j} = \boldsymbol{e}_{1+jk+i}, \quad j = 0, \ldots, J-1,\ i = 0, \ldots, k-1\,,$$

and the upper triangular matrix $\boldsymbol{R} \in \mathbb{C}^{m\times m}$ column-wise by

$$\boldsymbol{R}\,\boldsymbol{e}_{1+iJ+j} = \boldsymbol{T}^{j-1}\boldsymbol{e}_{1+iJ}, \quad j = 0, \ldots, J-1,\ i = 0, \ldots, k-1\,.$$

[4] meaning $\mathbf{V} = (\mathbf{M}^{-1}\mathbf{A})\mathbf{U}$, $\mathbf{V}^H\mathbf{M}\mathbf{V} = \mathbf{I}$.

Finally define $\tilde{\boldsymbol{U}} = \boldsymbol{U}_{:,1:J:m} \in \mathbb{C}^{N\times k}$. *Then* $\boldsymbol{R}$ *has full rank and the following equality holds:*

$$\boldsymbol{U}\boldsymbol{R} = K_J(\boldsymbol{M}^{-1}\boldsymbol{A}; \tilde{\boldsymbol{U}})\,\boldsymbol{\Pi}\,. \tag{2.4}$$

We call $(\tilde{\boldsymbol{U}}, \boldsymbol{\Pi}, \boldsymbol{R})$ *a* short representation *of* $\boldsymbol{U}$.

This result can be found by simple calculations. $\mathbf{R}$ has $\mathcal{O}(mJ)$ entries and can be computed in $\mathcal{O}(mJ)$. In [16, fig. 5 & 6] structures of $\mathbf{R}, \boldsymbol{\Pi}$ are given for sample values of k, J.

The proposed definition and lemma are useful in the sense, that instead of $\mathbf{U}$ only the matrices $\tilde{\mathbf{U}}, \mathbf{R}$ must be stored[5]. This short representation of $\mathbf{U}$ enables the computation of matrix-vector-products with $\mathbf{U}$ and its transpose by applying matrix-vector-products with $K_J(\mathbf{M}^{-1}\mathbf{A}; \tilde{\mathbf{U}})$ and its transpose, respectively. The matrix-vector-products with $K_J(\mathbf{M}^{-1}\mathbf{A}; \tilde{\mathbf{U}})$ and its transpose in turn can be computed by a Horner and a power scheme, respectively.

The idea for computation of a product of the form $\mathbf{z} = K_J(\mathbf{M}^{-1}\mathbf{A}; \tilde{\mathbf{U}}) \cdot \mathbf{y}$, where $\mathbf{y} = (\mathbf{y}_1^T, \mathbf{y}_2^T, \ldots, \mathbf{y}_J^T)^T \in \mathbb{C}^m$, $\mathbf{y}_i \in \mathbb{C}^k$, is given by the Horner-type scheme

$$\mathbf{z} = \tilde{\mathbf{U}}\mathbf{y}_1 + \mathbf{M}^{-1}\mathbf{A}\left(\tilde{\mathbf{U}}\mathbf{y}_2 + \mathbf{M}^{-1}\mathbf{A}\,(\ldots)\right)\,,$$

or, equivalently, algorithmically by:

1: $\mathbf{z} := \tilde{\mathbf{U}}\tilde{\mathbf{y}}_J$
2: **for** $j = J-1,\ J-2,\ \ldots,\ 1$ **do**
3: $\quad \mathbf{z} := \mathbf{A}\mathbf{z}$
4: $\quad \mathbf{z} := \mathbf{M}^{-1}\mathbf{z}$
5: $\quad \mathbf{z} := \mathbf{z} + \tilde{\mathbf{U}}\tilde{\mathbf{y}}_j$
6: **end for**

Similarly, a product $\mathbf{y} = K_J(\mathbf{M}^{-1}\mathbf{A}; \tilde{\mathbf{U}})^H \cdot \mathbf{z}$, where, as above, $\mathbf{y} = (\mathbf{y}_1^T, \mathbf{y}_2^T, \ldots, \mathbf{y}_J^T)^T$, can be computed by the power scheme

$$\mathbf{y}_j = \tilde{\mathbf{U}}^H\left((\mathbf{A}\mathbf{M}^{-1})^{j-1}\,\mathbf{z}\right)\,, \quad j = 1, \ldots, J\,,$$

or, equivalently, by:

1: **for** $j = 1, \ldots, J-1$ **do**
2: $\quad \tilde{\mathbf{y}}_j := \tilde{\mathbf{U}}^H\mathbf{z}$
3: $\quad \mathbf{z} := \mathbf{M}^{-1}\mathbf{z}$
4: $\quad \mathbf{z} := \mathbf{A}\mathbf{z}$
5: **end for**
6: $\tilde{\mathbf{y}}_J := \tilde{\mathbf{U}}^H\mathbf{z}$

The computational cost for applying $K_J(\mathbf{M}^{-1}\mathbf{A}; \tilde{\mathbf{U}})$ or its transpose to a vector is $J-1$ MV-s with $\mathbf{A}$ and $\mathbf{M}^{-1}$ and J matrix-vector-products with the dense $N \times k$-matrix $\tilde{\mathbf{U}}$.

Now the recycling strategy is as follows: We solve the first system with PCR. During this process we collect data $\tilde{\mathbf{U}}$, $\mathbf{T}$ to obtain a short representation of the search space $\mathcal{U}$ that PCR used. Then for a subsequent system with RHS $\mathbf{b}^{(\iota)}$ we can

[5] The permutation represented by $\boldsymbol{\Pi}$ is uniquely defined by k, J, cf. Lemma 1, hence the matrix $\boldsymbol{\Pi}$ does not have to be stored.

use the short representation to recycle that search space $\mathcal{U}$, i.e. by formula (2.3) we compute a residual-optimal solution in $\mathbf{x}_0^{(\iota)} + \mathcal{U}$. Such a solution from a recycled search space is called *recycling solution*. However, when using such a recycling strategy in practice, there are two main issues that have not been discussed yet:

- *Problem A:* Due to loss of orthogonality in $\mathbf{V}$ in equation (2.3) and stability issues with the block Krylov matrix and $\mathbf{R}$ in (2.4) the size of k and J is limited in practice (depending on the system). If, for example, the size is limited to $k = 10$ and $J = 6$, how to recycle search spaces of more than 60 dimensions?
- *Problem B:* When a recycling solution is found, how to improve it residual-optimally to satisfy a given accuracy demand?

In the next subsection we discuss one possible approach to overcome both problems.

2.3. Building Orthogonal Krylov Subspaces. We first provide a simple result and then discuss its benefit.

LEMMA 2 (KSS-Orthogonality). *From* (2.1) *consider the last columns* $\boldsymbol{d}_m, \boldsymbol{v}_m$ *of the basis matrices* $\boldsymbol{D}, \boldsymbol{V}$. *With these we define the modified system matrix*

$$\tilde{\boldsymbol{A}} := (\boldsymbol{I} - \boldsymbol{d}_m \boldsymbol{v}_m^H)\, \boldsymbol{A}\, \underline{(\boldsymbol{I} - \boldsymbol{v}_m \boldsymbol{d}_m^H)}\,. \tag{2.5}$$

Let $\boldsymbol{r} \in \mathbb{C}^N$ *be an arbitrary vector that satisfies* $\boldsymbol{V}^H \boldsymbol{r} = \boldsymbol{0}$. *Define its* KSS

$$\mathcal{W} = \boldsymbol{M}^{-1} \cdot \mathcal{K}_n(\tilde{\boldsymbol{A}} \boldsymbol{M}^{-1}; \boldsymbol{r})\,. \tag{2.6}$$

Note that we used the modified system matrix in (2.6). *Then the space* $\mathcal{W}$ *satisfies*

$$\boldsymbol{M}^{-1}\boldsymbol{A} \cdot \mathcal{W} \perp_{\boldsymbol{M}} \boldsymbol{M}^{-1}\boldsymbol{A} \cdot \mathit{range}(\boldsymbol{U})\,. \tag{2.7}$$

For $\mathbf{M} = \mathbf{I}$ this result is obvious. For $\mathbf{M} \neq \mathbf{I}$ it follows by simple substitution of a symmetric splitting of $\mathbf{M}$. As $\mathbf{r} \perp_{\mathbf{M}} \mathbf{v}_m$ by requirement, the underlined factor in (2.5) can be dropped, similar as in (1.5).

In the following we explain the usefulness of this result: Call $\mathbf{W}$ a basis of $\mathcal{W}$ and $\mathbf{Z}$ its orthogonal image. Further let $\mathbf{U}$ and its orthogonal image $\mathbf{V}$ be given by (2.1). Condition (2.7) means $\mathbf{V} \perp_{\mathbf{M}} \mathbf{Z}$, or $[\mathbf{V}, \mathbf{Z}]^H \mathbf{M} [\mathbf{V}, \mathbf{Z}] = \mathbf{I}$. Thus a residual-optimal solution in $\mathit{range}(\mathbf{U}) + \mathit{range}(\mathbf{W})$ can be computed by iteratively orthogonalizing[6] a residual onto columns of $[\mathbf{V}, \mathbf{Z}]$ and updating its numerical solution accordingly by columns of $[\mathbf{U}, \mathbf{W}]$.

Regarding *Problem A*, assume we used PCR to construct a large Lanczos decomposition with basis $\mathbf{U} \in \mathbb{C}^{N \times m}$ and tridiagonal matrix $\mathbf{T} \in \mathbb{C}^{m \times m}$. Let us further assume that due to the large size of $\mathbf{U}$ its numerical orthogonality has very low accuracy. So large errors in the recycling solution would occur if we applied (2.3) directly with $\mathbf{U}$. Additionally let us assume that $\mathbf{T}$ is ill-conditioned, so $\mathbf{R}$ of $\mathbf{U}$'s short representation would probably also be ill-conditioned.

To limit these influences on the recycling solution, we suggest to split the basis $\mathbf{U}$ into a few *basis blocks*, e.g. two. Then we have $\mathbf{U} = [\mathbf{U}_\mathrm{I}, \mathbf{U}_\mathrm{II}]$. The Lanczos decomposition has the property that neighboring basis vectors are orthogonal w.r.t. each other with high accuracy [18]. So as now the blocks $\mathbf{U}_i \in \mathbb{C}^{m_i \times m_i}$, $m_\mathrm{I} + m_\mathrm{II} =$

[6] To be more precise: For a residual $\mathbf{r}$ we orthogonalize the preconditioned residual $\hat{\mathbf{r}} \equiv \mathbf{M}^{-1}\mathbf{r}$ onto columns of $\mathbf{V}$ w.r.t. $\langle \cdot, \cdot \rangle_\mathbf{M}$. Note that the computation of the $\mathbf{M}$-scalar-product can be avoided using $\langle \mathbf{V}, \hat{\mathbf{r}} \rangle_\mathbf{M} \equiv \langle \mathbf{V}, \mathbf{r} \rangle$. In Algorithm 1 we used $\langle \mathbf{V}, \mathbf{r} \rangle \equiv \langle \mathbf{D}, \hat{\mathbf{r}} \rangle$ instead.

m, $i \in \{\mathrm{I}, \mathrm{II}\}$, are smaller than $\mathbf{U}$, their respective orthogonality property[7] is more accurate. With this result, we can use (2.3) to orthogonalize a residual subsequently on each block (comparable to as MINRES and PCR do for columns).

We did not address the question yet how to keep each basis block in memory. Our solution is to store a short representation for each block. To do so we split the tridiagonal matrix $\mathbf{T}$ from above into $\mathbf{T} = \mathrm{diag}(\mathbf{T}_\mathrm{I}, \mathbf{T}_\mathrm{II}) + \mathbf{E}$, where $\mathbf{T}_i \in \mathbb{C}^{m_i \times m_i}$, $i \in \{\mathrm{I}, \mathrm{II}\}$, are Hermitian tridiagonal matrices with probably better condition than $\mathbf{T}$. $\mathbf{E}$ consists only of $\beta_{m_\mathrm{I}+1}$ in two subdiagonal entries. With each $\mathbf{T}_i$ we can build a short representation for each $\mathbf{U}_i$ of higher precision. This is done by using (2.5), details can be found in section 2.3.1 below.

Regarding *Problem B*, we consider the situation that for a subsequent system a recycling solution $\mathbf{x}^{(\iota)}$ has been computed with the old recycling data from (2.1), such that $\mathbf{V}^H \mathbf{r}^{(\iota)} = \mathbf{0}$. To improve the solution, we take it as initial guess for PCR with an orthogonalization approach, where $\mathbf{A}$ is replaced by the modified system matrix $\tilde{\mathbf{A}}$ from (2.5), where m denotes the number of columns of $\mathbf{U}, \mathbf{D}, \mathbf{V}$ (i.e. in the formula we use the last column of each of these three matrices). As all conditions of Lemma 2 are satisfied, PCR will improve the solution in a residual-optimal way, as it builds a search space $\mathcal{W}$ that fulfills condition (2.7).

Implementation. To embed the orthogonalization approach for PCR in algorithm 1 to address *Problem B*, `A_func` in line 5 must be replaced by `modA_func`:

1: **procedure** MODA_FUNC($\mathbf{M}, \mathbf{A}, \hat{\mathbf{r}}$)
2: $\quad[\hat{\mathbf{v}}, \hat{\mathbf{d}}, \hat{\mathbf{u}}]$:=`A_func`$(\mathbf{M}, \mathbf{A}, \hat{\mathbf{r}})$ // $\hat{\mathbf{r}} \perp \mathbf{d}_m$ *already*
3: $\quad\gamma := \langle \mathbf{v}_m, \hat{\mathbf{d}} \rangle$
4: $\quad\hat{\mathbf{v}} := \hat{\mathbf{v}} - \gamma\, \mathbf{v}_m, \quad \hat{\mathbf{u}} := \hat{\mathbf{u}} - \gamma\, \mathbf{u}_m$
5: $\quad$**return** $\hat{\mathbf{v}}, \hat{\mathbf{d}}, \hat{\mathbf{u}}$
6: **end procedure**

We see that for the recycling the column $\mathbf{d}_m$ does not need to be stored.

2.3.1. An Explanatory Example. To illustrate the recycling strategy we present a fictitious example where $n_\mathrm{Eqns} = 3$ subsequent systems $\mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \mathbf{b}^{(3)}$ are solved with system matrix $\mathbf{A}$ and preconditioner $\mathbf{M}^{-1}$.

We start by solving $\mathbf{b}^{(1)}$ with PCR within $\hat{m} = 102$ iterations. During this we store the tridiagonal matrix $\mathbf{T}$ up to iteration $m = 90$ and from the full basis matrix $\mathbf{U}$ the columns $\mathbf{u}_1, \mathbf{u}_6, \mathbf{u}_{11}, \mathbf{u}_{16}, \mathbf{u}_{21}, \ldots, \mathbf{u}_{86}$ (choosing $J = 5$). In addition to that we also store the columns $\mathbf{u}_{30}, \mathbf{v}_{30}$, $\mathbf{u}_{60}, \mathbf{v}_{60}$, and $\mathbf{u}_{90}, \mathbf{v}_{90}$. From the stored columns we construct the three matrices $\tilde{\mathbf{U}}_\mathrm{I} = [\mathbf{u}_1, \mathbf{u}_6, \mathbf{u}_{11}, \mathbf{u}_{16}, \mathbf{u}_{21}, \mathbf{u}_{26}]$, $\tilde{\mathbf{U}}_\mathrm{II} = [\mathbf{u}_{31}, \mathbf{u}_{36}, \mathbf{u}_{41}, \mathbf{u}_{46}, \mathbf{u}_{51}, \mathbf{u}_{56}]$ and $\tilde{\mathbf{U}}_\mathrm{III} = [\mathbf{u}_{61}, \mathbf{u}_{66}, \mathbf{u}_{71}, \mathbf{u}_{76}, \mathbf{u}_{81}, \mathbf{u}_{86}]$. We declare the block matrices $\mathbf{U}_\mathrm{I} = [\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_{30}]$, $\mathbf{U}_\mathrm{II} = [\mathbf{u}_{31}, \mathbf{u}_{32}, \ldots, \mathbf{u}_{60}]$ and $\mathbf{U}_\mathrm{III} = [\mathbf{u}_{61}, \mathbf{u}_{62}, \ldots, \mathbf{u}_{90}]$ for later reference, but do not store them.

Choosing the parameters $k = 6, J = 5$, the corresponding index permutation is described by $\mathbf{\Pi} \in \mathbb{C}^{30\times 30}$ in Lemma 1. With $\mathbf{T} = \mathrm{diag}(\mathbf{T}_\mathrm{I}, \mathbf{T}_\mathrm{II}, \mathbf{T}_\mathrm{III}) + \mathbf{E}$, $\mathbf{T}_i \in \mathbb{C}^{30\times 30}$ the diagonal blocks of $\mathbf{T} \in \mathbb{C}^{90\times 90}$, we construct $\mathbf{R}_i \in \mathbb{C}^{30\times 30}$ from input arguments $k, J, \mathbf{T}_i$ for all $i \in \{\mathrm{I}, \mathrm{II}, \mathrm{III}\}$, as described in Lemma 1. We denote the modified system

[7] i.e. $\langle \mathbf{M}^{-1}\mathbf{A}\mathbf{U}_i\mathbf{e}_j, \mathbf{M}^{-1}\mathbf{A}\mathbf{U}_i\mathbf{e}_h \rangle_\mathbf{M} = \delta_{j,h}$, $\forall i, h = 1, \ldots, m_i$, $i \in \{\mathrm{I}, \mathrm{II}\}$.

matrices[8] by

$$\begin{aligned}\tilde{\mathbf{A}}_{\mathrm{I}} &:= (\mathbf{I} - \mathbf{d}_{30}\mathbf{v}_{30}^H)\,\mathbf{A},\\ \tilde{\mathbf{A}}_{\mathrm{II}} &:= (\mathbf{I} - \mathbf{d}_{60}\mathbf{v}_{60}^H)\,\mathbf{A},\\ \tilde{\mathbf{A}}_{\mathrm{III}} &:= (\mathbf{I} - \mathbf{d}_{90}\mathbf{v}_{90}^H)\,\mathbf{A}.\end{aligned}$$

With these we obtain the following short representations of $\mathbf{U}_{\mathrm{I}}, \mathbf{U}_{\mathrm{II}}, \mathbf{U}_{\mathrm{III}}$:

$$\begin{aligned}\mathbf{U}_{\mathrm{I}}\mathbf{R}_{\mathrm{I}} &= K_5(\mathbf{M}^{-1}\mathbf{A}; \tilde{\mathbf{U}}_{\mathrm{I}})\,\mathbf{\Pi},\\ \mathbf{U}_{\mathrm{II}}\mathbf{R}_{\mathrm{II}} &= K_5(\mathbf{M}^{-1}\tilde{\mathbf{A}}_{\mathrm{I}}; \tilde{\mathbf{U}}_{\mathrm{II}})\,\mathbf{\Pi},\\ \mathbf{U}_{\mathrm{III}}\mathbf{R}_{\mathrm{III}} &= K_5(\mathbf{M}^{-1}\tilde{\mathbf{A}}_{\mathrm{II}}; \tilde{\mathbf{U}}_{\mathrm{III}})\,\mathbf{\Pi}\,.\end{aligned}$$

Next, for the solution of $\mathbf{b}^{(2)}$ we recycle the first 90 dimensions of the former search basis $\mathbf{U}$ as follows: From the initial guess $\mathbf{x}_0^{(2)}$ we compute the associated residual $\mathbf{r}_0^{(2)}$ and then apply formula (2.3) for $\mathbf{U}_{\mathrm{I}}$ to derive the improved solution $\mathbf{x}_{\mathrm{I}}^{(2)}$. For $\mathbf{x}_{\mathrm{I}}^{(2)}$ in turn we compute the associated residual $\mathbf{r}_{\mathrm{I}}^{(2)}$ and then apply formula (2.3) for $\mathbf{U}_{\mathrm{II}}$ to derive the improved solution $\mathbf{x}_{\mathrm{II}}^{(2)}$. Finally for $\mathbf{x}_{\mathrm{II}}^{(2)}$ we compute the associated residual $\mathbf{r}_{\mathrm{II}}^{(2)}$ and then apply formula (2.3) for $\mathbf{U}_{\mathrm{III}}$ to obtain the recycling solution $\mathbf{x}_{\mathrm{III}}^{(2)}$ of the recycling basis $[\mathbf{U}_{\mathrm{I}}, \mathbf{U}_{\mathrm{II}}, \mathbf{U}_{\mathrm{III}}]$. For each block $2J$ MV-s are necessary for application of formula (2.3). Thus we needed 30 MV-s so far and storage for 24 column vectors of length N, but found a residual-optimal solution in a 90-dimensional space.

Now we want to perform a-posteriori iterations to improve $\mathbf{x}_{\mathrm{III}}^{(2)}$. For that purpose we run PCR on the system $\mathbf{A}\mathbf{x}^{(2)} = \mathbf{b}^{(2)}$ with initial guess $\mathbf{x}_{\mathrm{III}}^{(2)}$ and preconditioner $\mathbf{M}^{-1}$, where `A_func` in PCR is replaced by `modA_func` for $\tilde{\mathbf{A}}_{\mathrm{III}}$. After $m = 42$ iterations PCR has found a residual-optimal solution in the whole search space of $90 + 42 = 132$ dimensions with sufficiently small residual. The basis of this 42-dimensional search space is denoted by $\mathbf{U}_{\mathrm{IV}} = [\mathbf{u}_{91}, \mathbf{u}_{92}, \ldots, \mathbf{u}_{132}]$.

In total, to compute the solution for $\mathbf{b}^{(2)}$ residual-optimally in a 132-dimensional search space, only 24 column-vectors needed to be stored (except resources during PCR) and $30 + 42 = 72$ MV-s with $\mathbf{A}$ and $\mathbf{M}^{-1}$ needed to be computed.

The solution to system $\mathbf{b}^{(3)}$ can be computed in the same way as for system $\mathbf{b}^{(2)}$, by simply reusing the recycling basis blocks $[\mathbf{U}_{\mathrm{I}}, \mathbf{U}_{\mathrm{II}}, \mathbf{U}_{\mathrm{III}}]$ by short representations, respectively, by subsequently orthogonalizing a residual of a numerical solution w.r.t. each block. As a remark: Alternatively a short representation of $\mathbf{U}_{\mathrm{IV}}$ could be fetched during the a-posteriori iterations for system $\mathbf{b}^{(2)}$ and added to the other three recycled blocks. Then by virtue of (2.3) for system $\mathbf{b}^{(3)}$ a recycling solution $\mathbf{x}_{\mathrm{IV}}^{(3)}$ (with a residual that is orthogonal to the images of all four recycling basis blocks) can be computed. However for the a-posteriori-iterations to improve $\mathbf{x}_{\mathrm{IV}}^{(3)}$ in an optimal sense we need then the modified system matrix

$$\tilde{\mathbf{A}}_{\mathrm{IV}} = (\mathbf{I} - [\mathbf{d}_{90}, \mathbf{d}_{132}][\mathbf{v}_{90}, \mathbf{v}_{132}]^H)\,\mathbf{A}\,\underline{(\mathbf{I} - [\mathbf{v}_{90}, \mathbf{v}_{132}][\mathbf{d}_{90}, \mathbf{d}_{132}]^H)}\,,$$

because $\mathbf{u}_{90}$ did not originate from the KSS where $\mathbf{u}_{132}$ originated from. (The underlined factor can be dropped.)

[8] We stress that the vectors of $\mathbf{D}$, that are used in these matrices, do not have to be stored, as we do not apply the modified system matrices explicitly but in a chain with the preconditioner, cf. `modA_func`.

2.4. Computational Cost. For simplicity we assume that we recycle a KSS of only one system (i.e. we neglect the remark above), by ℓ basis blocks of identical width $m = kJ$, each of which is shortly represented by a dense $N \times k$-matrix. (For the above example we had $\ell = 3$, $k = 6$, $J = 5$.) We neglect storage and computational cost in $\mathcal{O}(\ell m J)$ (with $J \leqslant m$).

To orthogonalize a residual w.r.t. one block, one needs to compute $2J$ MV-s with both $\mathbf{A}$ and $\mathbf{M}^{-1}$, $2J$ scalar-products, $6J$ AXPYs and J matrix-vector-products with both the dense $N \times k$-matrix and its transpose.

In total, for recycling a $\ell k J$-dimensional search space, $\ell(k+2)$ column-vectors need to be stored. The $+2$ occurs due to the two vectors $\mathbf{u}_{m_i}, \mathbf{v}_{m_i}$ that must be additionally stored for each block $i = 1, \ldots, \ell$ for the modified system matrix.

To construct a residual-optimal recycling solution in the recycled $\ell k J$-dimensional search space, the cost is dominated by $2\ell J$ MV-s with $\mathbf{A}$ and $\mathbf{M}^{-1}$, and ℓJ matrix-vector-products with a dense $N \times k$-matrix and its transpose, respectively.

3. Numerical Experiments. In this section we present four test cases in which we compare the proposed method SR-PCR-ap to the well-known method MINRES [19], as it is implemented in Matlab. All test cases are from partial differential equations (PDE) discretizations. The first two are taken from the Florida Sparse Matrix Collection [22], the second two stem from finite element discretizations.

The "Right" Right-Hand Sides. In [27, sec. 3.7] the author states that the choice of RHS-es is not crucial for a test case. We disagree in that point as for matrices representing discrete elliptic differential operators (which are of practical relevance) it is well-known that for high-frequent solutions a KSS solver needs more iterations to resolve these frequencies.

It also seems to be a common practice to define the RHS by the (non-smooth) image of a smooth solution vector, e.g. in [23]. However, when solving in [23, sec. 6.4 (c)] for the RHS u (smooth) instead of F, the according solution vector has high frequencies (which is physically meaningful, cf. to solutions of Turing bifurcation problems). To solve for this RHS u, GMRES and MINRES would need twice as many iterations as for F. To sum it up, choosing a RHS from a smooth solution vector may provide convergence curves that differ much from those for a RHS coming from practical applications.

In addition to that, symmetry both of the domain and the RHS-function can influence the convergence speed extremely. As an example, for $-\Delta u = 1$, $u|_{\partial\Omega} = 0$ on a Cartesian isotropic grid of the domain $\Omega = (0,1)^2$, one can expect at least $\sqrt{8}$ times faster convergence than for a random RHS due to symmetry (which of course also affects round-off errors and thus stability, and due to lower dimension of the full KSS also the superlinear convergence).

To avoid any confusion about initial guesses and whether it is useful to choose the initial guess of the next system as the solution to the former system, the initial guess $\mathbf{x}_0 = \mathbf{0}$ is chosen for each RHS and ten RHS-es $\mathbf{b}^{(1)}, \ldots, \mathbf{b}^{(10)}$ are constructed for each test case, that are orthogonal w.r.t. each other.

For our test cases there is additionally the question how to choose a series of RHS-es: When e.g. choosing two RHS-es $\mathbf{b}^{(1)}$, $\mathbf{b}^{(2)}$, where $\mathbf{b}^{(2)}$ is orthogonal to the image of the former search space[9] then any KSS recycling method is by construction just as good as its respective method without recycling, cf. example 1. So we have to ensure that the recycled search space of $\mathbf{b}^{(1)}$ contains good solution candidates

[9] i.e. $\mathbf{b}^{(2)} \perp_{\mathbf{M}} \mathbf{A}\mathbf{M}^{-1} \cdot \mathcal{K}_m(\mathbf{A}\mathbf{M}^{-1}; \mathbf{b}^{(1)})$

for the subsequent RHS-es (*utility*) and on the other hand we have to ask which kind of sequences occur in practice (*practical relevance*). We make three suggestions, going from practical relevance to utility. For each sequence the RHS-es are chosen orthogonal w.r.t. each other and independent of n_{Eqns}, the number of RHS-es, using an orthogonal[10] basis of a KSS to some *starting vector* $\mathbf{d}$:

1. *Sequence A:* When the next RHS depends on a linear combination of former solutions, then the choice

$$span\{\mathbf{b}^{(1)}, \ldots, \mathbf{b}^{(q)}\} = \mathcal{K}_q(\mathbf{A}^{-1}; \mathbf{d}) \quad \forall\, q = 1, \ldots, n_{\text{Eqns}} \tag{3.1}$$

   seems reasonable. For Hermitian positive definite systems we know by the convergence results from [27, eqn. 6.107],[25, chap. 5.2], that $\mathcal{K}_m(\mathbf{A}; \mathbf{b}^{(1)})$ contains good solution candidates for the subsequent RHS-es of this sequence for sufficiently large m. However, this is not clear for the search space with preconditioning $\mathbf{M}^{-1} \cdot \mathcal{K}_m(\mathbf{A}\mathbf{M}^{-1}; \mathbf{b}^{(1)})$.
2. *Sequence B:* We generalize (3.1) under preconditioning.

$$span\{\mathbf{b}^{(1)}, \ldots, \mathbf{b}^{(q)}\} = \mathcal{K}_q(\mathbf{M}\mathbf{A}^{-1}; \mathbf{M}^{-1}\mathbf{d}) \quad \forall\, q = 1, \ldots, n_{\text{Eqns}} \tag{3.2}$$

   results in (3.1), if $\mathbf{M}$ is splitted (e.g. Cholesky) and substituted into $\mathbf{A}$. For the same reasons as above it is known that $\mathbf{M}^{-1} \cdot \mathcal{K}_m(\mathbf{A}\mathbf{M}^{-1}; \mathbf{b}^{(1)})$ contains good solution candidates for the subsequent RHS-es of this sequence for sufficiently large m.
3. *Sequence C:* This last sequence is defined for systems where strong preconditioning is required to be able to find an iterative solution at all. Numerical experiments indicate that the nearer the preconditioner is to $\mathbf{A}^{-1}$, the more suited the recycling space is *only* to $\mathbf{b}^{(1)}$, cf. results for sequence A in section 3.3.1. Therefor we consider the following test sequence, where the influence of the preconditioner on the utility is attenuated:

$$span\{\mathbf{b}^{(1)}, \ldots, \mathbf{b}^{(q)}\} = \mathcal{K}_q(\mathbf{A}\mathbf{M}^{-1}; \mathbf{d}) \quad \forall\, q = 1, \ldots, n_{\text{Eqns}}\,. \tag{3.3}$$

   The search space $\mathbf{M}^{-1} \cdot \mathcal{K}_m(\mathbf{A}\mathbf{M}^{-1}; \mathbf{b}^{(1)})$ often provides good solution candidates for these RHS-es, but it does not contain the corresponding exact solution. In this setting recycling could be probably replaced by parallel solves for the different RHS-es, as there is no dependence of a RHS on a former solution (as no $\mathbf{A}^{-1}$ occurs).

Example 1 (Limits of Recycling). *We stress that there exist sequences of linear systems for which recycling is completely useless. As an example, let* $N \in 2\mathbb{N}$, $\boldsymbol{A} = 1/(N+1)^2 \cdot tridiag(-1, 2, -1) \in \mathbb{R}^{N \times N}$, $\boldsymbol{M} = I$, *and consider a sequence of* 2 *RHS-es* $\boldsymbol{b}^{(1)} = (\boldsymbol{1}^T, \boldsymbol{1}^T)^T$, $\boldsymbol{b}^{(2)} = (-\boldsymbol{1}^T, \boldsymbol{1}^T)^T$ *with* $\boldsymbol{1} \in \mathbb{R}^{N/2}$. *The equations can be interpreted as discretizations of Poisson problems* $-\Delta u = f^{(\iota)}$ *in* $\Omega = (-1, 1)$, $u(-1) = u(1) = 0$ *with RHS-functions* $f^{(1)}(x) = 1$, $f^{(2)}(x) = sign(x)$.

Obviously for the first system the RHS, the solution as well as the vectors $\boldsymbol{A}^j \boldsymbol{b}^{(1)}$ $\forall j \in \mathbb{N}$ *are symmetric w.r.t. the symmetry axis* $x = 0$. *In contrast, for the second system the RHS, the solution and the vectors* $\boldsymbol{A}^j \boldsymbol{b}^{(2)}$ $\forall j \in \mathbb{N}$ *are antisymmetric w.r.t.*

[10] w.r.t. the Euclidian scalar-product

$x = 0$. *Consequently, the recycling spaces* $\mathcal{U}^{(\iota)} := \boldsymbol{M}^{-1} \cdot \mathcal{K}_N(\boldsymbol{A}\boldsymbol{M}^{-1}; \boldsymbol{b}^{(\iota)})$, $\iota = 1, 2$, *have the following properties:*

$$\mathcal{U}^{(2)} \perp_{\boldsymbol{M}} \mathcal{U}^{(1)}\,,$$
$$\boldsymbol{M}^{-1}\boldsymbol{A} \cdot \mathcal{U}^{(2)} \perp_{\boldsymbol{M}} \boldsymbol{M}^{-1}\boldsymbol{A} \cdot \mathcal{U}^{(1)}\,,$$

which means that 1.) the recycling space $\mathcal{U}^{(1)}$ *is perpendicular to* $\boldsymbol{x}^{(2)}$ *and 2.) its image is orthogonal to* $\boldsymbol{b}^{(2)}$. *Thus for the second system a recycling solution in* $\mathcal{U}^{(1)}$ *does not improve the initial guess* $\boldsymbol{x}_0 = \boldsymbol{0}$, *regardless of whether a residual-optimal or an error-optimal recycling method would be used.*

3.1. CURLCURL0. In this test case [26] only the indefinite symmetric matrix $\mathbf{A} \in \mathbb{R}^{11083\times 11083}$ is given. For sake of simplicity we chose $\mathbf{d} = \mathbf{A}\,\mathbf{1}$.

System without preconditioning. We first study the unpreconditioned system (i.e., $\mathbf{M} = \mathbf{I}$) with system matrix $\mathbf{A}$ and the RHS $\mathbf{b} = \mathbf{d}$. The convergence of $\|\mathbf{r}\|_2/\|\mathbf{b}\|_2$ for MINRES is shown in black in figure 3.1.

Next, we want to estimate an appropriate block size for recycling, cf. section 2.3. Therefore the matrix $\mathbf{Q} := \mathbf{V}^H\mathbf{M}\,\mathbf{V}$ is computed with $\mathbf{V}$ from (2.1) of the solution process with PCR. $\mathbf{Q}$ should be the identity but due to round-off errors it also has non-zero off-diagonal entries. This affects the stability of the Gram-Schmidt orthogonalization that is hidden in (2.3). The $\log_{10}|\cdot|$ for each entry of $\mathbf{Q}$ is shown in the left of figure 3.2. Right in the figure $\log_{10}|\cdot|$ is given for each entry of a symmetric matrix $\mathbf{G}$ that has the entries $g_{i,j} = \kappa_2(\mathbf{T}_{i:j,i:j})$, i.e. the 2-condition of a section of the tridiagonal matrix from (2.1). We discussed in section 2.3 that the ill-conditioning of these sub-matrices affects the accuracy of the short representations. So as a rule of thumb, the bandwidth in which the entries of both matrices of figure 3.2 are small limits the block size for each basis block. These two plots give an estimate how large each basis block can be chosen. Taking larger blocks is beneficial as increasing the number of blocks leads to an increased memory consumption.

From $\mathbf{Q}$ in fig. 3.2 we see that e.g. a block of the first 50 basis columns could be recycled (as these are nearly orthogonal w.r.t. each other), then maybe a small block from columns 51 to 60 could be recycled, and then a block from columns 60 to 100, and so on. However this seems complicated and not so efficient, as we would need many small blocks, which needs more storage and computational effort than a few large blocks. We now turn to a study of the block size for the preconditioned system.

Preconditioned System. We will see that by use of preconditioning the orthogonality and conditioning properties are improved, i.e., the off-diagonal elements of $\mathbf{Q}$ and $\mathbf{G}$ become smaller. Here the preconditioner

$$\mathbf{M} = \mathit{sign}(\mathrm{diag}(\mathbf{A})) \cdot \mathit{tridiag}(\mathbf{A}) \tag{3.4}$$

is used which is Hermitian positive definite. The convergence of $\|\mathbf{r}\|_2/\|\mathbf{b}\|_2$ for PMINRES (i.e. preconditioned MINRES with $\mathbf{M}$) is given in fig. 3.1 in thick gray. The iterates are algebraically equivalent to those of PCR, but as PCR does not iteratively compute $\mathbf{r}$, we plot $\|\mathbf{M}^{-1}\mathbf{r}\|_2/\|\mathbf{M}^{-1}\mathbf{b}\|_2$ instead (dashed curve in fig. 3.1). In the remainder we assume both residual measures to be comparable.

Extracting the basis columns from the solution process of PCR, the matrices $\mathbf{Q}, \mathbf{G}$ are shown in fig. 3.3. Comparing the result to fig. 3.2, we see that now the basis blocks can be chosen larger, e.g. the first 150 columns are nearly orthogonal w.r.t. each other. However, the conditioning of $\mathbf{T}$'s sub-blocks still limits the block sizes. As for practical problems the matrices $\mathbf{Q}, \mathbf{G}$ cannot be computed, we will make very defensive choices for the block sizes in the following.

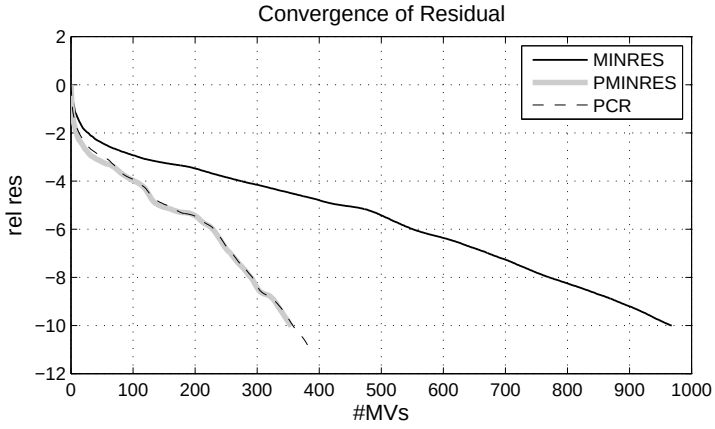

FIGURE 3.1. *Convergence of* MINRES,PMINRES *and* PCR *for CURLCURL0 with* RHS $\boldsymbol{b} = \boldsymbol{A}\,\mathbf{1}$.

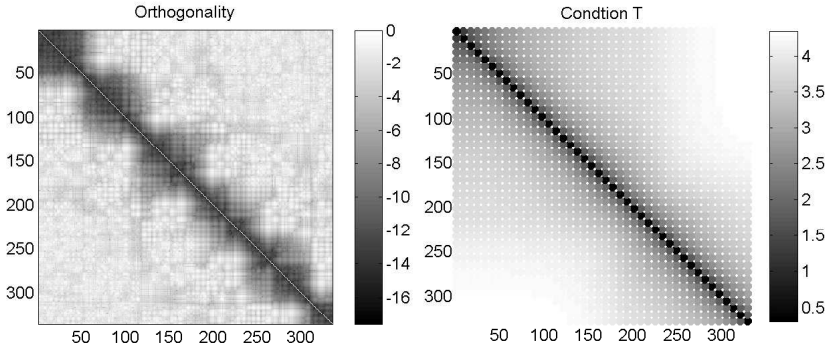

FIGURE 3.2. *Matrices (left)* $\boldsymbol{Q}$ *and (right)* $\boldsymbol{G}$ *for unpreconditioned CURLCURL0.*

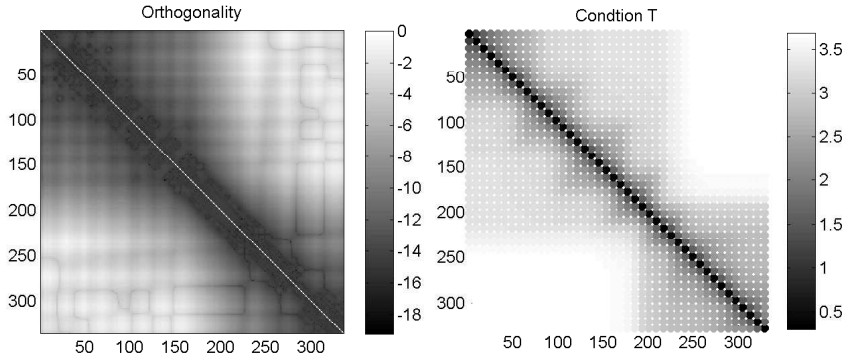

FIGURE 3.3. *Matrices (left)* $\boldsymbol{Q}$ *and (right)* $\boldsymbol{G}$ *for preconditioned CURLCURL0.*

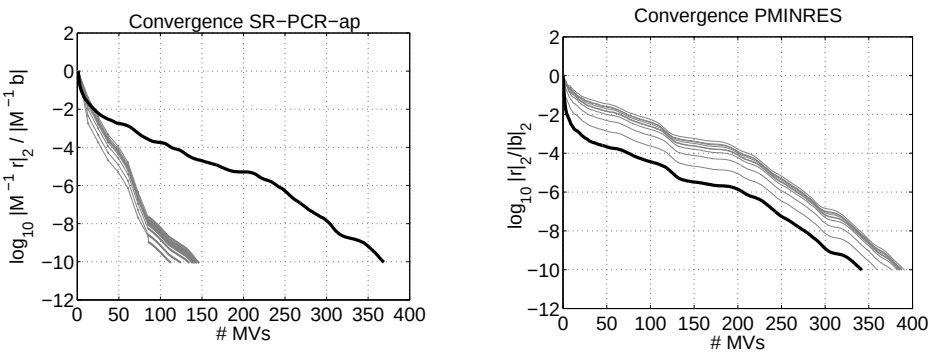

(a) (b)

FIGURE 3.4. *Convergence for all RHS-es of sequence C of CURLCURL0 for (a) SR-PCR-ap, (b) PMINRES; thick black = first RHS, thin gray = subsequent RHS-es.*

Solving Sequence C. For $\mathbf{M}, \mathbf{A}, \mathbf{d}$ we construct $\mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \ldots, \mathbf{b}^{(10)}$ from sequence C. To solve the sequence, we first solve for $\mathbf{b}^{(1)}$ with PCR and fetch short representations for its basis blocks. For this we choose $\ell = 7$ subsequent basis blocks of each $k = 8$ vectors and $J = 6$ (yielding a block size of $kJ = 48$ for each block). By this the search space $\mathbf{M}^{-1} \cdot \mathcal{K}_{336}(\mathbf{A}\mathbf{M}^{-1}; \mathbf{b}^{(1)})$ can be recycled with a cost of 70 stored columns and 84 MV-s. Afterwards we solve for $\mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \ldots, \mathbf{b}^{(10)}$ with SR-PCR-ap, i.e. we use formula (2.3) subsequently for each block, followed by residual-optimal post-iterations.

Fig. 3.4 shows the convergence of SR-PCR-ap (left) in comparison to PMINRES (right). One can observe that recycling leads to a dramatic acceleration of convergence speed in this example. E.g., for a required accuracy of 10^{-8} SR-PCR-ap needs only 1/3 of the MV-s compared to PMINRES.

3.2. SHERMAN1. This test set [24] provides a symmetric indefinite matrix $\mathbf{A} \in \mathbb{R}^{1000\times 1000}$, $\kappa_2(\mathbf{A}) \approx 2.3 \cdot 10^4$, with one vector that we use for $\mathbf{d}$, the starting vector of the RHS-sequence.

We directly go for the preconditioned system using a preconditioner $\mathbf{M}$ as in (3.4) for CURLCURL0. For the RHS-es we used sequence B.

As for CURLCURL0, we solve for $\mathbf{b}^{(1)}$ with PCR, collect short representations, and then recycle these basis information using the proposed method SR-PCR-ap. As parameters we choose $\ell = 2$ blocks of each $k = 8$ stored columns and $J = 7$, and thus recycle the search space $\mathbf{M}^{-1} \cdot \mathcal{K}_{112}(\mathbf{A}\mathbf{M}^{-1}; \mathbf{b}^{(1)})$ storing 20 columns and applying 28 MV-s.

Fig. 3.5 shows a comparison of the convergence of the residuals for SR-PCR-ap with those of PMINRES. We see that by use of the recycling information the proposed method converges faster for each RHS, but for later RHS-es the number of MV-s grows, as the recycling space becomes out of date.

3.3. Finite Element Test Cases. The two following test cases stem from FE discretizations of two different PDE-s in 2D, namely a Poisson and Stokes problem. The FE discretizations are carried out using SimpleFEM [17]. The two-dimensional domain Ω and the corresponding triangular mesh are shown in figure 3.6 (a). The Neumann boundary part Γ_N is given in thick gray (inner boundary; hole), the Dirichlet

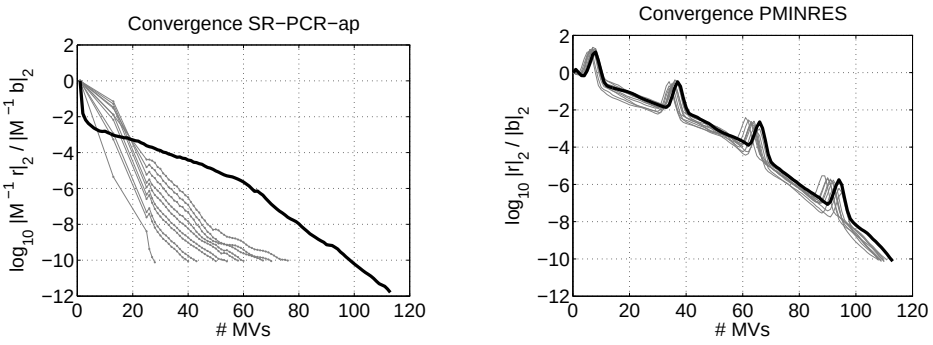

(a) (b)

FIGURE 3.5. *Convergence for all RHS-es of sequence B of SHERMAN1 for (a) SR-PCR-ap, (b) PMINRES; thick black = first RHS, thin gray = subsequent RHS-es.*

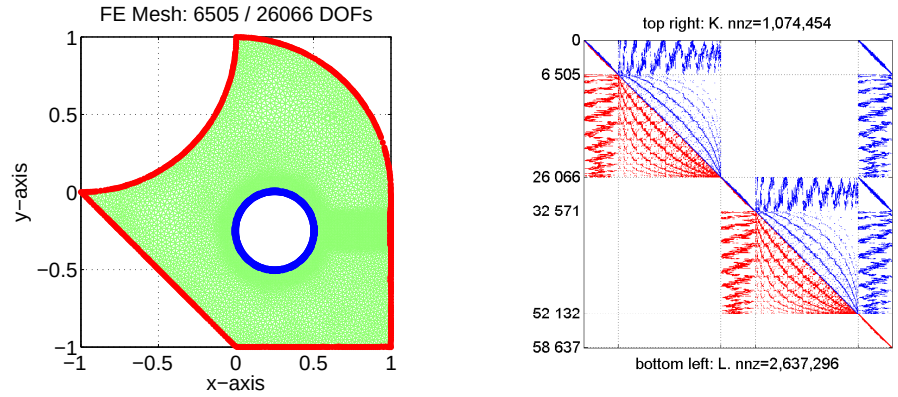

(a) (b)

FIGURE 3.6. *(a) Finite Element Mesh with nodes: red = Dirichlet boundary nodes, blue = Neumann boundary nodes (inner hole). (b) Pattern of Stokes Matrix (blue, only upper right part) and its Preconditioner (red).*

part Γ_D in black. On this mesh, piecewise quadratic FE are chosen for the discretization of the Poisson problem. For the Stokes problem piecewise linear and piecewise quadratic FE are used for the discretization of the pressure and the velocity field, respectively.

3.3.1. Poisson Problem. We compute a numerical solution ϕ to

$$\begin{aligned} -\Delta\phi &= 0 \quad \text{in } \Omega \\ \vec{\nabla}\phi \cdot \vec{n} &= 0 \quad \text{on } \Gamma_N \\ \phi &= g_D \quad \text{on } \Gamma_D \qquad \text{with } g_D(x,y) := \begin{cases} 1, & \text{if } x = 1 \\ 0, & \text{else} \end{cases} \end{aligned}$$

using piecewise quadratic FE on the mesh shown in fig. 3.6 (a). The resulting stiffness matrix $\mathbf{A} \in \mathbb{R}^{26066\times 26066}$ is symmetric positive definite with $\kappa_1(\mathbf{A}) \approx 6.5 \cdot 10^4$.

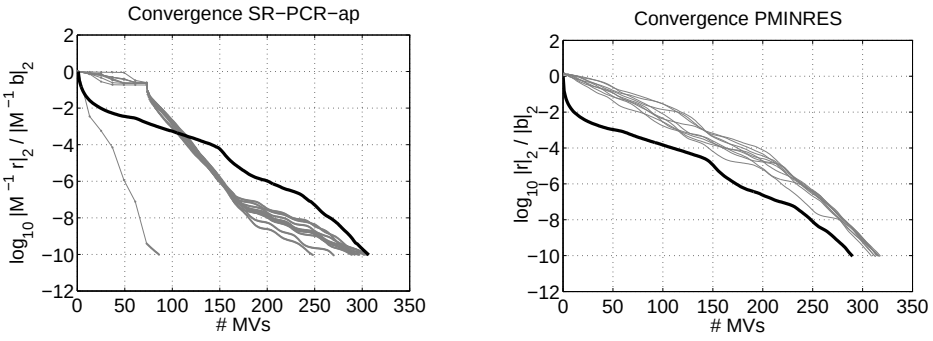

(a) (b)

FIGURE 3.7. *Convergence for all RHS-es of sequence A of POISSON for (a) SR-PCR-ap, (b) PMINRES; thick black = first RHS, thin gray = subsequent RHS-es.*

The non-zero RHS $\mathbf{d}$ contains the contributions from the non-homogeneous Dirichlet boundary data. We choose the preconditioner $\mathbf{M} = \mathbf{L}\mathbf{L}^H$, where $\mathbf{L}$ is the incomplete Cholesky decomposition [27, Algo. 10.7] of $\mathbf{A}$.

We try $\ell = 6$ blocks, each of dimension 48, by choosing $k = 8$, $J = 6$ for each block. By this we recycle the search space $\mathbf{M}^{-1} \cdot \mathcal{K}_{288}(\mathbf{A}\mathbf{M}^{-1}; \mathbf{b}^{(1)})$ by storing 60 columns and performing 72 MV-s.

Solving Sequence A. We first solve for sequence A of $\mathbf{d}$, cf. (3.1). For these RHS-es the convergence curves of SR-PCR-ap and MINRES are compared in figure 3.7. We observe that SR-PCR-ap has a faster convergence after the recycling phase, possibly due to the space dimensions on which the residuals are orthogonalized. Anyway, for the first RHS (from which the recycling space is built) the convergence speed-up by recycling is much different from that of the other RHS-es.

A possible reason for this might be that due to preconditioning the recycled search space and the RHS-es do not fit together well. As already noted for sequence A, one can only expect that the the recycling space without preconditioning $\mathcal{K}_m(\mathbf{A}; \mathbf{d})$ (for m sufficiently large) contains good candidates for subsequent RHS-es. Instead, we used the recycling space $\mathbf{M}^{-1}\,\mathcal{K}_m(\mathbf{M}^{-1}\,\mathbf{A}; \mathbf{d})$.

Solving Sequence B. As a second test we solve for sequence B of $\mathbf{d}$, cf. (3.2). For this sequence we obtain completely different convergence curves than for sequence A, cf. figure 3.8. This is due to the fact that now the recycled search space contains useful candidates for subsequent RHS-es for the preconditioned case. This can be directly seen in figure 3.8 as the convergence curves do not differ very much from each other.

From the comparison of these two sequences we draw the following conclusion: *If in a practical use case a sequence of type B with an arbitrary but known matrix $\boldsymbol{M}$ occurs (such as e.g. a mass matrix from FE discretizations), then for sake of a useful recycling space it might be advantageous to choose a preconditioner that is somehow "similar" to $\boldsymbol{M}$.* This statement implies sequences of type A, as for these $\mathbf{M} = \mathbf{I}$ holds. So to a certain degree there is a trade-off between choosing an efficient versus a "similar" preconditioner when combining it with recycling.

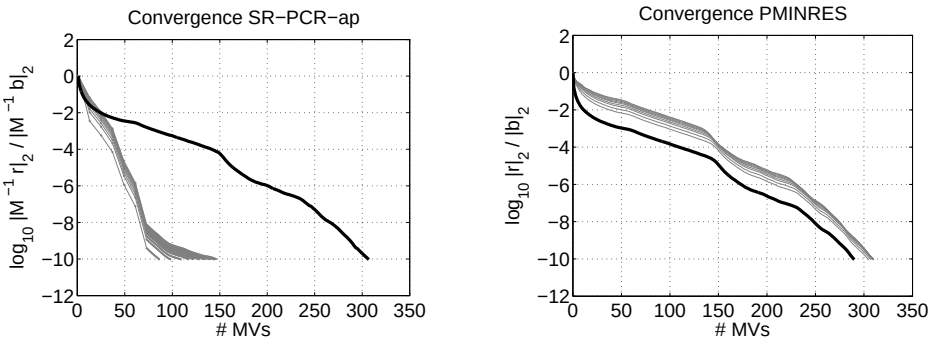

(a) (b)

FIGURE 3.8. *Convergence for all RHS-es of sequence B of POISSON for (a) SR-PCR-ap, (b) PMINRES; thick black = first RHS, thin gray = subsequent RHS-es.*

3.3.2. Stokes Problem. We consider the following Stokes problem

$$\begin{array}{rrrll} -\Delta u & & + \partial_x p & = 0 & \text{in } \Omega \\ & -\Delta v & + \partial_y p & = 0 & \text{in } \Omega \\ \partial_x u & + \partial_y v & & = 0 & \text{in } \Omega \\ n_x \nabla u & + n_y \nabla v & - p\vec{n} & = 0 & \text{on } \Gamma_N \\ & & u & = 0 & \text{on } \Gamma_D \\ & & v & = g_D & \text{on } \Gamma_D \end{array}$$

with p, u, v denoting the pressure, velocity in x and y direction, respectively. With the chosen boundary conditions this models flow in a lid-driven cavity.

Discretization. We compute a numerical solution u_h, v_h, p_h with u_h, v_h piecewise quadratic FE functions and p_h piecewise linear FE functions on the mesh in fig. 3.6 (a). By taking the discrete Laplacian $\mathbf{A}$ from the Poisson problem above, $\mathbf{M}_p \in \mathbb{R}^{6505\times 6505}$ the FE mass matrix for linear shape functions (index $_p$ indicates pressure space), and $\mathbf{B}_i = \langle \psi, \partial_i \varphi \rangle_{L^2(\Omega)}$, $i \in \{x, y\}$, for quadratic shape functions φ as columns and linear shape functions ψ as rows, we solve

$$\mathbf{K}\,\mathbf{x} \equiv \begin{bmatrix} \mathbf{A} & \mathbf{O} & \mathbf{B}_x^H \\ \mathbf{O} & \mathbf{A} & \mathbf{B}_y^H \\ \mathbf{B}_x & \mathbf{B}_y & \mathbf{O} \end{bmatrix} \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \\ -\mathbf{p} \end{pmatrix} = \mathbf{g}\,.$$

For the symmetric indefinite saddle point matrix $\mathbf{K} \in \mathbb{R}^{58637\times 58637}$ with $\kappa_1(\mathbf{K}) \approx 3.8 \cdot 10^6$, we choose the preconditioner $\mathbf{M} = \mathbf{L}\mathbf{L}^H$ with $\mathbf{L} = \mathrm{diag}(\mathbf{L}_A, \mathbf{L}_A, \mathbf{L}_{M_p})$ where $\mathbf{L}_A$ and $\mathbf{L}_{M_p}$ are incomplete Cholesky decompositions of $\mathbf{A}$ and $\mathbf{M}_p$, respectively, using Matlab's `ichol` with threshold $\epsilon = 10^{-4}$. The structure of $\mathbf{K}$ and $\mathbf{L}$ is given in fig. 3.6 (b). The numerical solution for the RHS $\mathbf{g} = (\mathbf{0}^T, \mathbf{d}^T, \mathbf{0}^T)^T$ with $\mathbf{d}$ as in the Poisson problem is shown in fig. 3.9.

Solving Sequences B and C. We keep $\ell = 2$ blocks of each $k = 10$ vectors and $J = 4$ to recycle the search space $\mathbf{M}^{-1}\cdot\mathcal{K}_{80}(\mathbf{A}\mathbf{M}^{-1}; \mathbf{b}^{(1)})$. We emphasize that for the Stokes problem we use a strong preconditioner. To investigate its influence on the utility of the recycling space, we compare the results for sequence B and C with starting vector $\mathbf{g}$, respectively. Note that by construction of both sequences the computed flow field is not discretely divergence-free as the third sub-vector of the RHS-es is inconsistent.

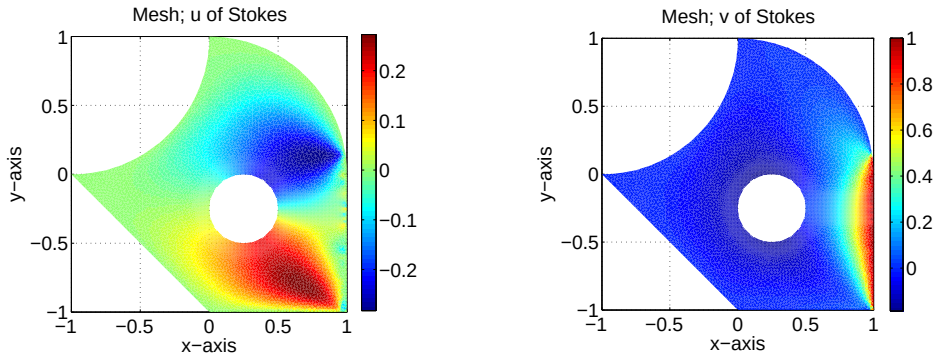

(a) (b)

FIGURE 3.9. *Flow field of the numerical solution of the Stokes problem (a) in x- and (b) in y-direction.*

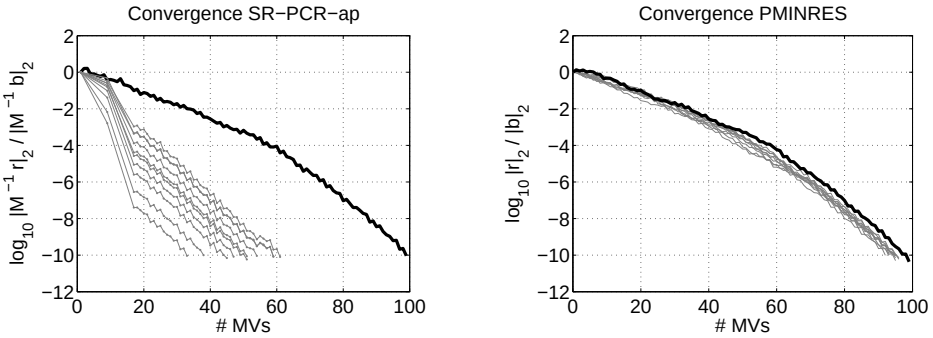

(a) (b)

FIGURE 3.10. *Convergence for all RHS-es of sequence B of STOKES for (a) SR-PCR-ap, (b) PMINRES; thick black = first RHS, thin gray = subsequent RHS-es.*

The convergence curves for each sequence are given in fig. 3.10 and fig. 3.11, respectively. From the figures we see that SR-PCR-ap offers a reduction in the average number of MV-s of about 50% for both sequences compared to PMINRES. However, for sequence B the convergence curves of SR-PCR-ap differ more from each other, with growing iteration numbers for later RHS-es. This indicates that the recycling space gradually becomes out of date and that this effect is more pronounced for sequence B than for sequence C.

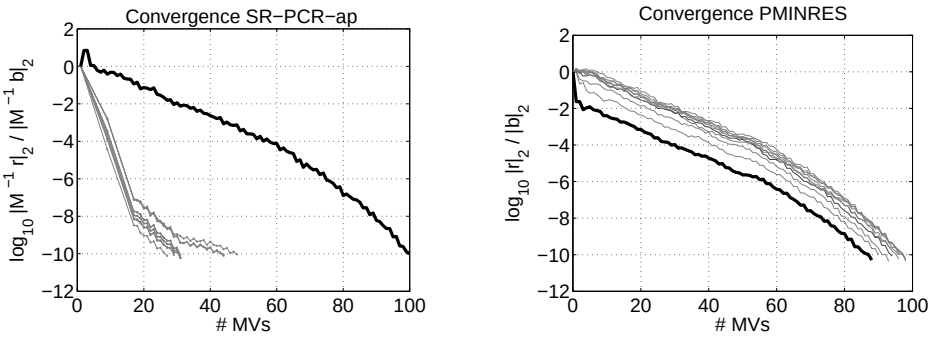

(a) (b)

FIGURE 3.11. *Convergence for all RHS-es of sequence C of STOKES for (a) SR-PCR-ap, (b) PMINRES; thick black = first RHS, thin gray = subsequent RHS-es.*

4. Conclusions and Outlook. The numerical experiments indicate that despite the delicate numerical properties of power and Horner schemes our method is quite practicable and efficient. In the experiments the number of stored columns could be reduced to $1/4^{\text{th}}$ of the dimension of the recycling space. The recycling approach shown here for PCR can be easily adapted to other Krylov subspace methods. From the experiments (cf. figs. 3.2 and 3.3) we deduce that the proneness to rounding errors of SR-PCR-ap is correlated to the conditioning of the respective system.

Depending on the type of sequence, the recycling strategy turns out to be more or less successful, especially when combined with a preconditioner, cf. figs. 3.7 and 3.8. As the preconditioner plays a dominant role in the construction of the recycling space it crucially influences its utility for subsequent systems. So in the design of a preconditioner the type of the RHS-sequence might be taken into consideration when combining preconditioning with recycling. This is a current topic of on-going research.

REFERENCES

[1] S. Wang and E. de Sturler and G. H. Paulino, Large-scale topology optimization using preconditioned Krylov subspace methods with recycling, Int. J. for Num. Meth. in Engineering, Vol. 69(12), pp. 2441-2468, 2006.

[2] K. Ahuja and E. de Sturler and S. Gugercin and E. R. Chang, Recycling BiCG with an Application to Model Reduction, SIAM J. Sci. Comput. Vol. 34, No. 4, pp. A1925-A1949, 2012.

[3] K. Ahuja and E. de Sturler and P. Benner, Recycling BiCGSTAB with an Application to Parametric Model Order Reduction, MPI Magdeburg preprints, pp. 13-21, 2013.

[4] P. Benner and L. Feng, Recycling Krylov Subspaces for Solving Linear Systems with successively changing Right-Hand-Sides arising in Model Reduction, Lecture Notes in Electrical Engineering, Vol. 74, pp. 125-140, Springer 2011.

[5] M. Parks and E. de Sturler and G. Mackey and D.D. Johnson and S. Maiti, Recycling Krylov subspaces for sequences of linear systems, SIAM J. Sci. Comput. Vol. 28(5), pp. 1651-1674, 2006.

[6] J. Bolten and N. Bozovic and A. Frommer, Preconditioning of Krylov subspace methods using recycling in Lattice QCD computations, Proc. Appl. Math. Mech., Vol. 13, pp. 413-414, 2013.

[7] K. Mohamed and S. Nadarajah and M. Paraschivoiu, Krylov Recycling techniques for unsteady simulation of turbulent aerodynamic flows, 26th international congress of the aeronautical sciences, 2008.

[8] Z. Ye and Z. Zhu and J. R. Phillips, Generalized Krylov Recycling Methods for Solution of Multiple Related Linear Equation Systems in Electromagnetic Analysis, Design Automation Conference 2008, p. 682-687.

[9] K. M. Soodhalter and D. E. Szyld and F. Xue, Krylov Subspace Recycling for Sequences of Shifted Linear Systems, Elsvier J. Appl. Num. Math. Vol. 81, pp. 105-118, 2014.

[10] R. B. Morgan, A restarted GMRES method augmented with eigenvectors. SIAM J. Matrix Anal. Appl., 16:1154-1171, 1995.

[11] H. A. van der Vorst and C. Vuik, GMRESR: A family of nested GMRES methods. Num. Lin. Alg. with Appl., 1:369-386, 1994.

[12] R. B. Morgan, Implicitly restarted GMRES and Arnoldi methods for nonsymmetric systems of equations. SIAM J. Matrix Anal. Appl., 21 (2000), pp. 1112-1135.

[13] M. Habu and T. Nodera, GMRESM algorithm with changing the restart cycle adaptively, Proceedings of ALGORITHMY 2000, Conference on Scientific Computing, pp. 254-263.

[14] E. de Sturler, Truncation Strategies for optimal Krylov subspace methods, SIAM J. Numer. Anal., Vol. 36(3), pp. 864-889, 1999.

[15] R. B. Morgan, GMRES with Deflated Restarting, SIAM J. Sci. Comput., 24(1), pp. 20-37, 2002.

[16] M. P. Neuenhofen, Short-Recurrence and -Storage Recycling of large Krylov-Subspaces for Sequences of Linear Systems with changing Right-Hand-Sides, Technical Report, available on arXiv: 1512.05101, 2015.

[17] M. P. Neuenhofen, Simple Finite Element Code, `MartinNeuenhofen.de/Simple_FEM/SimpleFEM.html`, 2015.

[18] H. D. Simon, The Lanczos Algorithm With Partial Reorthogonalization, Mathematics of Computation, 46(165):115-142, 1984.

[19] C. C. Paige and M. A. Saunders, Solution of sparse indefinite Systems of linear equations, SIAM J. Numer. Anal. Vol. 12(4), pp. 617-629, 1975.

[20] S. T. Choi, C. C. Paige and M. A. Saunders, MINRES-QLP: A Krylov Subspace Method for Indefinite or Singular Symmetric Systems, SIAM J. Sci. Comput., 33(4), pp. 1810-1836, 2011.

[21] K. Hayami, On the Behaviour of the Conjugate Residual Method for Singular Systems, Technical Report, National Institute of Informatics, Tokyo 2011.

[22] T. A. Davis and Y. Hu, The University of Florida Sparse Matrix Collection. ACM Transactions on Mathematical Software, 38(1):1-25. `http://www.cise.ufl.edu/research/sparse/matrices`. 2011.

[23] G. L. G. Sleijpen and M. B. van Gijzen, Exploiting BiCGstab(ℓ) Strategies to Induce Dimension Reduction, SIAM J. Sci. Comput. 32(5):2687-2709, 2010.

[24] Duff, I. S. and R. G. Grimes and J. G. Lewis, Sparse Matrix Problems. ACM Trans. on Mathematical Software, 14(1):1-14, 1989.

[25] S. Gross and A. Reusken, Numerical Methods for Two-phase Incompressible Flows. First edition 2011, Springer Series in Computational Mathematics, Vol. 40, ISBN 978-3-642-19685-0.

[26] Z. Bai, Krylov subspace techniques for reduced-order modeling of large-scale dynamical systems. Applied Numerical Mathematics, 43(1-2):9-44, 2002.

[27] Y. Saad, Iterative Methods for Sparse Linear Systems, 2nd edition, SIAM, 2000.